Mon. Not. R. Astron. Soc. **000**, 1–11 (2015) Printed 29 June 2016 (MN LATEX style file v2.2)

SPAMCART: a code for smoothed particle Monte Carlo radiative transfer

O. Lomax⋆, & A. P. Whitworth
School of Physics and Astronomy, Cardiff University, Cardiff CF24 3AA, UK

arXiv:1603.01223v4 [astro-ph.IM] 28 Jun 2016

29 June 2016

ABSTRACT

We present a code for generating synthetic SEDs and intensity maps from Smoothed Particle Hydrodynamics simulation snapshots. The code is based on the Lucy (1999) Monte Carlo Radiative Transfer method, i.e. it follows discrete luminosity packets as they propagate through a density field, and then uses their trajectories to compute the radiative equilibrium temperature of the ambient dust. The sources can be extended and/or embedded, and discrete and/or diffuse. The density is not mapped onto a grid, and therefore the calculation is performed at exactly the same resolution as the hydrodynamics. We present two example calculations using this method. First, we demonstrate that the code strictly adheres to Kirchhoff's law of radiation. Second, we present synthetic intensity maps and spectra of an embedded protostellar multiple system. The algorithm uses data structures that are already constructed for other purposes in modern particle codes. It is therefore relatively simple to implement.

Key words: radiative transfer – hydrodynamics – methods: numerical – ISM: dust

1 INTRODUCTION

Monte Carlo Radiative Transfer (MCRT) is a stochastic method for simulating radiative transfer through a medium. Individual MCRT calculations are accurate but slow. Fortunately, they are trivial to parallelise and therefore well suited to modern multi-threaded CPUs. MCRT is often used to post-process simulation snapshots and can be adapted to solve a variety of radiative processes. For example, codes such as HYPERION (Robitaille 2011) and RADMC-3D (Dullemond 2012) calculate the dust and molecular line emissivities. These data can then be used to generate realistic synthetic observations. Other MCRT codes include MOCASSIN (Ercolano et al. 2003), which models photoionisation fronts and emission line intensities from HII regions, and TORUS (Harries 2011), which models time-dependent radiative transfer.

Smoothed particle hydrodynamics (SPH) (Lucy 1977; Gingold & Monaghan 1977) is a mesh-free method of solving the equations of fluid dynamics. A discrete set of particles are used to model the density distribution. These particles are smoothed over one another using a kernel function with smoothing length h. By letting h vary with density, the particles can model a field spanning many orders of magnitude in density. This property makes SPH an attractive scheme for modelling astrophysical systems such as star forming clouds (e.g. Bate 2009; Lomax et al. 2014, 2015), galactic discs (e.g. Dobbs, Bonnell & Pringle 2006; Dobbs, Burkert & Pringle 2011) and the cosmic web (e.g. Springel & Hernquist 2003). Open-source SPH codes include GADGET-2 (Springel 2005) and GANDALF (Hubber & Rosotti 2013).

Performing MCRT calculations on density fields from SPH simulations usually involves mapping the particles onto an octree (e.g. Stamatellos & Whitworth 2005; Robitaille 2011; Rundle et al. 2010). This fundamentally changes the structure of the density field by adding noise and/or merging several particles into a single cell (we provide a brief analysis of this noise in Appendix A). Another option is to use the particle positions to construct a Voronoi tessellation (e.g. Hubber, Ercolano & Dale 2016). This is less noisy than an octree, although the implementation is more complicated. Nevertheless, both methods alter the SPH density structure by replacing smoothed particles with uniform density cells.

In this paper we present the new code SPAMCART (short for Smoothed PArticle Monte CArlo Radiative Transfer). The code performs MCRT calculations using the properties of the particles' smoothing kernels. Therefore, post-processing calculations are performed at *exactly* the same resolution as the SPH simulation. The algorithm is similar to other ray-tracing and Monte Carlo methods used in SPH, (e.g. Altay, Croft & Pelupessy 2008; Forgan & Rice 2010). However, this is the first time such an algorithm has been used to perform full MCRT calculations. The method utilises the pre-existing neighbour-finding/gravity

⋆ E-mail: oliver.lomax@astro.cf.ac.uk

tree found in modern SPH codes, and is therefore relatively simple to implement.

In §2 we describe the general Monte Carlo method we use for radiative transfer. In §3 we explain how this method can be applied to smoothed particles. In §4 we present some example calculations, including a physical benchmark. In §5 we list some future implementations and in §6 we present the summary.

2 MONTE CARLO RADIATIVE TRANSFER

The method presented here is an adaptation of the Lucy (1999) algorithm, modified to operate on smoothed particles rather than a grid of uniform density cells. Most of the operations involve drawing a random variate $\mathcal{U}[a,b]$ from the Uniform distribution in the interval (a,b).

2.1 Luminosity packet trajectories

A radiation source with luminosity L emits N_γ luminosity packets. Each packet is emitted from an origin $\boldsymbol{o}$ in a direction $\boldsymbol{n}$. If the radiator is an isotropic point source, $\boldsymbol{n}$ is given by polar angles

$$\begin{aligned} \phi &= \mathcal{U}[-\pi,\pi]\,, \\ \cos\theta_{\rm ISO} &= \mathcal{U}[-1,1]\,. \end{aligned} \tag{2.1}$$

If there is an external isotropic radiation field, $\boldsymbol{o}$ is a random point on a closed convex surface and $\boldsymbol{n}$ is given by polar angles

$$\begin{aligned} \phi &= \mathcal{U}[-\pi,\pi]\,, \\ \cos\theta_{\rm SUR} &= \sqrt{\mathcal{U}[0,1]}\,, \end{aligned} \tag{2.2}$$

where $\theta_{\rm SUR}$ is the angle between $\boldsymbol{n}$ and the inward surface normal. Each packet has energy $\varepsilon_0 \equiv (L\,\Delta t/N_\gamma)$[1] and wavelength λ, drawn randomly from the spectral energy distribution of the source. For the surface of a star, approximated as a blackbody, the equation

$$\frac{\int_0^\lambda B_{\lambda'}(T)\,{\rm d}\lambda'}{\int_0^\infty B_{\lambda'}(T)\,{\rm d}\lambda'} = \mathcal{U}[0,1] \tag{2.3}$$

is solved for λ, where T is the temperature of the source.

An individual packet travels an optical depth

$$\tau = -\ln(\mathcal{U}[0,1])\,. \tag{2.4}$$

This corresponds to a distance l, where

$$\tau = \int_0^l \rho(l')\,\chi_\lambda(l')\,{\rm d}l'\,. \tag{2.5}$$

Here, ρ is the density of the medium and χ_λ is the mass extinction coefficient. Solving Eqn. 2.5 for l is one of the more computationally demanding aspects of MCRT and is covered in the next section.

[1] If we assume radiative equilibrium, then the value of Δt is arbitrary and does not affect the calculation.

Once a packet has travelled distance l, it is either absorbed or scattered at position

$$\boldsymbol{o}_{\rm NEW} = \boldsymbol{x} + l\,\boldsymbol{o}\,. \tag{2.6}$$

A packet is absorbed if

$$a_\lambda < \mathcal{U}[0,1]\,, \tag{2.7}$$

where a_λ is the albedo of the medium. Otherwise the packet is scattered. An absorbed packet is immediately re-emitted with a new wavelength drawn from the emissivity distribution of the medium. This is achieved by solving the equation

$$\frac{\int_0^\lambda \kappa_{\lambda'}\,B_{\lambda'}(T)\,{\rm d}\lambda'}{\int_0^\infty \kappa_{\lambda'}\,B_{\lambda'}(T)\,{\rm d}\lambda'} = \mathcal{U}[0,1] \tag{2.8}$$

for λ, where $\kappa_\lambda \equiv (1-a_\lambda)\,\chi_\lambda$ is the monochromatic mass absorption coefficient and T is the temperature of the medium at $\boldsymbol{o}_{\rm NEW}$. If the packet is scattered, λ remains unchanged and $\boldsymbol{n}_{\rm NEW}$ is scattered away from $\boldsymbol{n}$ by a random angle $\theta_{\rm SCT}$. Here, $\theta_{\rm SCT}$ is drawn from the Henyey & Greenstein (1941) phase function:

$$\Phi(\theta) = \frac{1}{4\,\pi}\,\frac{1-g_\lambda^2}{[1+g_\lambda^2-2\,g_\lambda\,\cos(\theta)]^{3/2}}\,, \tag{2.9}$$

where g_λ is the mean scattering cosine. Hence, relative to the pre-scattering direction,

$$\begin{aligned} \phi &= \mathcal{U}[-\pi,\pi]\,, \\ \cos\theta_{\rm SCT} &= \frac{1}{2\,g}\left\{1+g^2-\left(\frac{1-g^2}{1+g\,\mathcal{U}[-1,1]}\right)^2\right\}\,. \end{aligned} \tag{2.10}$$

The process of absorption then re-emission and/or scattering is repeated until the packet exits the system.

2.2 Emissivity distribution

For a medium in radiative equilibrium, the energy emission rate is equal to the energy absorption rate $\dot{A}$, i.e.

$$\dot{A} = 4\,\kappa_{\rm P}\,\sigma_{\rm SB}\,T^4\,. \tag{2.11}$$

Here, $\sigma_{\rm SB}$ is the Stefan-Boltzmann constant, T is the local temperature and $\kappa_{\rm P}$ is the Planck mean absorption coefficient:

$$\kappa_{\rm P} = \frac{\int_0^\infty \kappa_\lambda\,B_\lambda(T)\,{\rm d}\lambda}{\int_0^\infty B_\lambda(T)\,{\rm d}\lambda}\,. \tag{2.12}$$

The average value of $\dot{A}$ within a volume V can be estimated by summing over the path lengths l_j of luminosity packets which pass through V,

$$\dot{A} \approx \frac{\varepsilon_0}{\Delta t}\,\frac{1}{V}\sum_j \kappa_{\lambda_j} l_j\,. \tag{2.13}$$

We can estimate T, and therefore the mass emissivity distribution, by calculating $\dot{A}$ and solving Eqn. 2.11 for T. However, Eqn. 2.13 depends on the trajectories of the luminosity packets and the trajectories depend on the emissivity distribution. Therefore several iterations are required to reach an equilibrium solution.

3 SPAMCART ALGORITHM

The main algorithm of the SPAMCART code consists of two tasks: (i) to calculate the trajectories of luminosity packets through an ensemble of smoothed particles and (ii) to estimate $\dot{A}$ for each particle. Here, we detail how these tasks are performed.

3.1 Density

For an ensemble of smoothed particles, each with position $\boldsymbol{x}_i$, mass m_i and smoothing length h_i, the density at any point in space is given by

$$\rho(\boldsymbol{x}) = \sum_i \frac{m_i}{h_i^3}\, w(s_i)\,, \quad s_i = \frac{\|\boldsymbol{x} - \boldsymbol{x}_i\|}{h_i}\,, \tag{3.1}$$

where $w(s)$ is the kernel function. Most kernel functions have compact support, i.e. they are only finite within ξ smoothing lengths. We use the M4 cubic spline kernel (Monaghan & Lattanzio 1985) which has $\xi = 2$ (see Appendix B). The smoothing length of each particle is calculated so that

$$\begin{aligned} h_i &= \eta \left(\frac{m_i}{\rho_i}\right)^{\frac{1}{3}}\,, \\ \rho_i &= \sum_j \frac{m_j}{h_i^3}\, w(s_{ij})\,, \quad s_{ij} = \frac{\|\boldsymbol{x}_j - \boldsymbol{x}_i\|}{h_i}\,. \end{aligned} \tag{3.2}$$

These equations are solved by iteration and we normally adopt $\eta = 1.2$ (as suggested by Price & Monaghan 2004).

3.2 Scatter Calculation

The value of some arbitrary quantity $Z(\boldsymbol{x})$ at an arbitrary position $\boldsymbol{x}$ can be estimated by *scattering* the same quantity, from each particle to $\boldsymbol{x}$, via the kernel function:

$$Z(\boldsymbol{x}) = \sum_i \frac{m_i}{\rho_i} \frac{Z_i}{h_i^3}\, w(s_i)\,. \tag{3.3}$$

The gradient $\nabla Z(\boldsymbol{x})$ may also be estimated from the gradient of the kernel function:

$$\begin{aligned} \nabla Z(\boldsymbol{x}) &= \sum_i \frac{m_i}{\rho_i} \frac{Z_i}{h_i^4}\, \nabla w(s_i)\,, \\ \nabla w(s_i) &= \frac{\boldsymbol{x} - \boldsymbol{x}_i}{\|\boldsymbol{x} - \boldsymbol{x}_i\|} \frac{\mathrm{d}}{\mathrm{d}s} w(s_i)\,. \end{aligned} \tag{3.4}$$

The gradient of the M4 kernel function is given in Appendix B.

3.3 Optical depth

The set of all points along a ray is defined by $\boldsymbol{x'} = \boldsymbol{o} + t\,\boldsymbol{n}$, where $0 \le t \le l$. Here, $\boldsymbol{o}$ is the origin, $\boldsymbol{n}$ is the direction unit vector and l is the length of the ray. The optical depth at wavelength λ along the ray is given by

$$\tau'(\boldsymbol{o}, \boldsymbol{n}, l) = \sum_i^{\text{RAY}} \chi_{\lambda i}\, \varsigma_i\,, \tag{3.5}$$

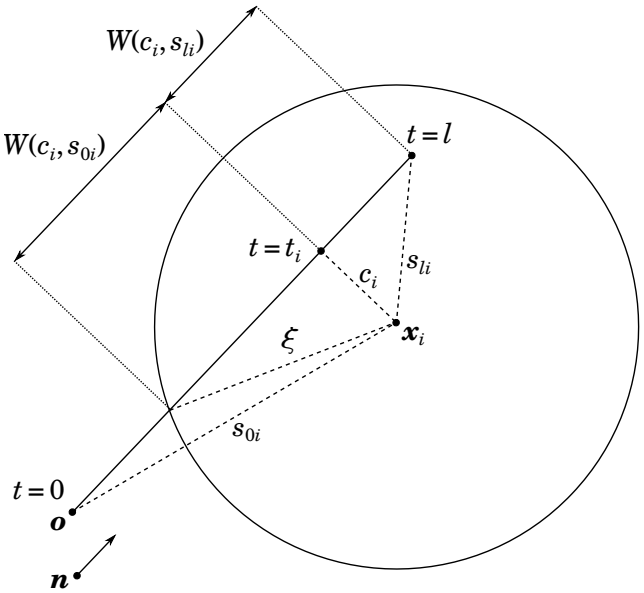

Figure 1. A diagram of line with origin $\boldsymbol{o}$, direction $\boldsymbol{n}$ and length l intersecting a particle at position $\boldsymbol{x_i}$ with compact support $s \le \xi$. The other terms are defined in §3.3 .

where

$$\begin{aligned} \varsigma_i &= \frac{m_i}{h_i^2} \begin{cases} W(c_i, s_{0i}) + W(c_i, s_{li}), & 0 \le t_i \le l; \\ |W(c_i, s_{0i}) - W(c_i, s_{li})|, & \text{otherwise}\,, \end{cases} \\ t_i &= (\boldsymbol{x}_i - \boldsymbol{o}) \cdot \boldsymbol{n}\,, \\ c_i &= \min\left(\frac{\|(\boldsymbol{o} + t_i\,\boldsymbol{n}) - \boldsymbol{x}_i\|}{h_i}, \xi\right)\,, \\ s_{0i} &= \min\left(\frac{\|\boldsymbol{o} - \boldsymbol{x}_i\|}{h_i}, \xi\right)\,, \\ s_{li} &= \min\left(\frac{\|\boldsymbol{o} + l\boldsymbol{n} - \boldsymbol{x}_i\|}{h_i}, \xi\right)\,, \\ W(c, s) &= \int_c^s \frac{w(s')\, s'}{\sqrt{s'^2 - c^2}}\, \mathrm{d}s'\,, \end{aligned} \tag{3.6}$$

and $\chi_{\lambda i}$ is the mass extinction coefficient of particle i. A diagram of this system for a single particle is given in Fig. 1 . The analytical form of $W(c, s)$ for the M4 kernel in given in Appendix B.

In order to calculate Eqn. 3.5, we must first identify all the particles which are intersected by the ray (see Fig. 2). The sum of all particle column density contributions is then used to calculate the total optical depth. Particle-ray intersections can be identified efficiently by walking a tree-structure (see Fig. 3) and opening cells which pass the slab test (e.g. Williams et al. 2005, see Appendix C). These trees are a standard element of SPH codes, used to optimise neighbour-finding and gravity calculations.

3.4 Propagating luminosity packets

We propagate a luminosity packet, starting at $\boldsymbol{o}$, travelling in direction $\boldsymbol{n}$. Initially, the ray has length l_0. Here l_0 is a modest overestimate of the intended value of l (see Eqn 2.5). A good choice of l_0 is provided by a second order Taylor

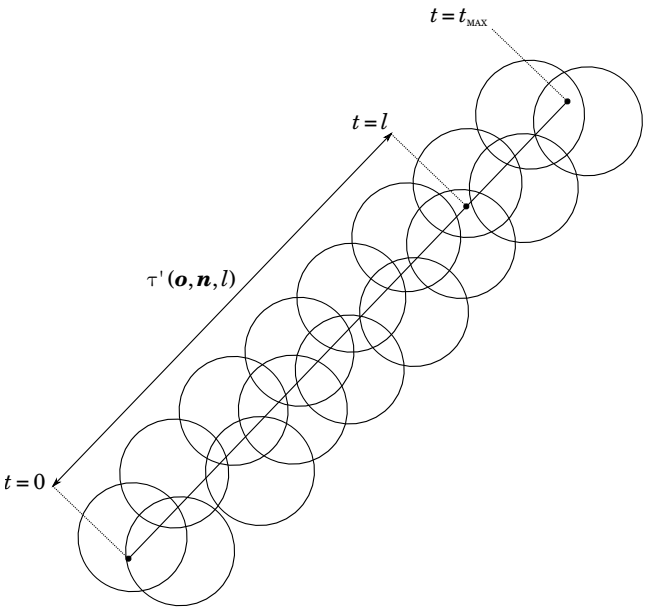

Figure 2. A diagram of all particles which are intersected by a ray of length $t_{\rm MAX}$. The optical depth along the ray may be calculated using Eqn. 3.5 . Note that $t_{\rm MAX}$ is greater than l. This is because l is calculated by iteration and the first iterate must be an overestimate. Also, because of the construction of Eqn. 3.5, particle column density beyond l does not contribute to the total optical depth.

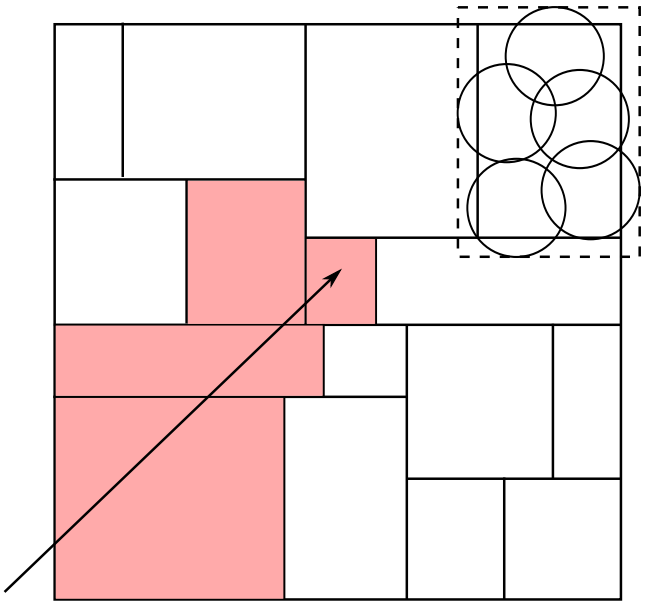

Figure 3. A schematic of the cells of a particle tree. Each cell has an axis-aligned bounding box (AABB) which encompasses the smoothing volume of all the particles in the cell (see Appendix C for more details). If the ray intersects the AABB of a cell with sub-cells, the sub-cells are recursively checked. If the cell has no sub-cells (shaded pink), all of its particles are added to the ray.

expansion at $l = 0$, i.e.

$$l_0 = \frac{2\,\zeta\,\tau}{l_\gamma^{-1}(\boldsymbol{o}) + \sqrt{\Delta}}\,, \qquad \Delta = \max([l_\gamma^{-1}(\boldsymbol{o})]^2 + 2\,\tau\,\nabla l_\gamma^{-1}(\boldsymbol{o})\cdot\boldsymbol{n}, 0)\,. \tag{3.7}$$

Here, the inverse mean-free-path $l_\gamma^{-1} = \chi_\lambda\,\rho$ is estimated via a scatter calculation and ζ is an overestimation factor which we set to 1.2 .

We use a modified Newton-Raphson root-finding method to iteratively solve $\tau' - \tau = 0$ for l (see Appendix D). We terminate the iterations when $|\tau' - \tau|/\tau \leq 0.01$, usually within four or five iterations. For some values of l_0, τ' is less than τ and no solution exists. Here, either the luminosity packet has left the ensemble of SPH particles or l_0 is too short. In the first case, the calculation for that luminosity packet is complete. In the second case, a new ray is constructed with the same direction and new origin $\boldsymbol{o}_{\rm NEW} = \boldsymbol{o} + l_0\,\boldsymbol{n}$. The calculation is then repeated with $\tau_{\rm NEW} = \tau - \tau'$.

3.5 Estimating the energy absorption rate

The energy absorption rate for an individual particle is estimated by summing the column densities along all packet trajectories:

$$\dot{A} = \frac{\varepsilon_0}{\Delta t}\frac{1}{m}\sum_j \kappa_{\lambda j}\,\varsigma_j\,. \tag{3.8}$$

Here ς_j is the column density along the path of luminosity packet j through the particle. Estimating $\dot{A}$ requires almost no computational expense as the values of ς_j have already been computed whilst propagating the luminosity packets.

The absorption rate from the previous iteration is used when randomly generating a new emission wavelength. The value of $\dot{A}$ at an arbitrary position can be found via a scatter calculation. The radiative equilibrium calculation is complete when the change in $\dot{A}_i$ (or temperature T_i) between iterations is less than a desired tolerance.

3.6 Generating intensity maps

We generate intensity maps of the dust emission via ray tracing. We construct a virtual rectangular screen facing the ensemble of particles. This screen is divided into $N_{\rm PIX}$ pixels. We construct $N_{\rm PIX}$ rays, each of which has infinite length and passes through the centre of a pixel with direction normal to the screen. A schematic of a single pixel is show in Fig. 4 . All intensity maps presented here have $N_{\rm PIX} = 200\times 200$.

Each ray intersects $N_{\rm RAY}$ particles. These are sorted into descending order of distance from the pixel. The intensity of the pixel is given by $I_{N_{\rm RAY}}$. This is calculated by iteration:

$$I_i = I_{i-1}\,\exp(-\chi_{\lambda i}\,\varsigma_i) + \frac{B_\lambda(T_i)\,\kappa_{\lambda i}}{\chi_{\lambda i}}[1 - \exp(-\chi_{\lambda i}\,\varsigma_i)]\,, \tag{3.9}$$

where $i = 1, 2, \ldots, N_{\rm PART}$. In the case where $i = 1$, I_0 is the background intensity.

The intensity from point sources (e.g. stars) may also be added to the map. This is achieved by (i) locating the pixel in which the source lies, (ii) calculating the optical depth τ between the source and the pixel and (iii)

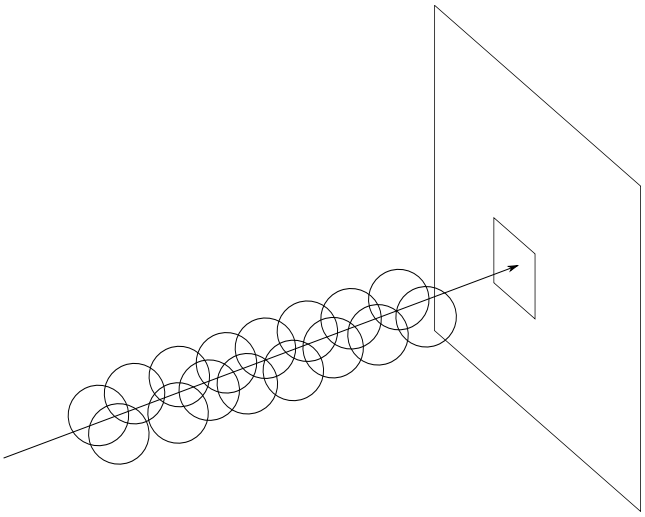

Figure 4. Schematic of a single pixel (small square) of an intensity map (large square). The ray passes through the centre of the pixel at an angle normal to the surface of the map.

adding $L_\lambda \, \mathrm{e}^{-\tau}/(4\pi \, A_{\mathrm{PIX}})$ to the pixel intensity. Here L_λ is the monochromatic luminosity of the source and A_{PIX} is the area of the pixel.

4 EXAMPLE CALCULATIONS

In the following calculations, we use the dust optical properties derived by (Li & Draine 2001). Here, the dust is a mixture of carbonaceous and amorphous silicate grains. The size distribution follows Weingartner & Draine (2001) with $R_v = 5.5$ and is normalised to give a dust-to-gas mass ratio of one percent. The code does not currently support imaging of scattered light (see §5), so we set the albedo to zero, i.e. $\kappa_\lambda \equiv \chi_\lambda$.

We present two examples. The first is a benchmark which tests the ability of the algorithm to reach thermal equilibrium with a radiation field. In the second example, we use the outputs of an SPH simulation to generate spectra and intensity maps of an embedded multiple system.

4.1 Blackbody radiation field

4.1.1 Undiluted blackbody radiation field

In an undiluted blackbody radiation field, the intensity $I_\lambda(\boldsymbol{x}, \boldsymbol{n})$ is equal to the Planck function $B_\lambda(T_{\mathrm{BB}})$. By construction, an absorbing/emitting object in this field has a surface intensity equal to $B_\lambda(T_{\mathrm{BB}})$. Following Kirchhoff's law of radiation—a good absorber is an equally good emitter—the surface temperature must therefore be equal to T_{BB}. Furthermore, because the emission spectrum is identical to the absorption spectrum, the object is invisible with respect to the background intensity.

4.1.2 Set up

We set up a sphere of $N_{\mathrm{PART}} = 3 \times 10^5$ particles. The sphere has radius $R = 3000\,\mathrm{au}$ and the density profile of a critical Bonnor-Ebert sphere, truncated at dimensionless radius $\xi = 6.451$. In the first instance, we give the sphere mass $M = 1\,\mathrm{M}_\odot$ (optically thin), in the second $M = 100\,\mathrm{M}_\odot$ (optically thick). We place the sphere in an undiluted $T_{\mathrm{BB}} = 60K$ blackbody radiation field. This is simulated by constructing a virtual shell which encompasses all the particles and their smoothing kernels, i.e. $R_{\mathrm{FIELD}} \gtrsim 2000\,\mathrm{au}$. The shell is given luminosity $L_{\mathrm{FIELD}} = 4\,\pi\,R_{\mathrm{FIELD}}^2\,\sigma_{\mathrm{SB}}\,T^4$. This luminosity is divided into $N_\gamma = 10^6$ luminosity packets, which are directed inwards towards the particles. The wavelength of each packet is drawn randomly from the Planck function $B_\lambda(T_{\mathrm{BB}})$. Each particle is given an initial temperature $T = 10\,\mathrm{K}$ and SPAMCART is iterated until the mean change in temperature is less than 1 K.

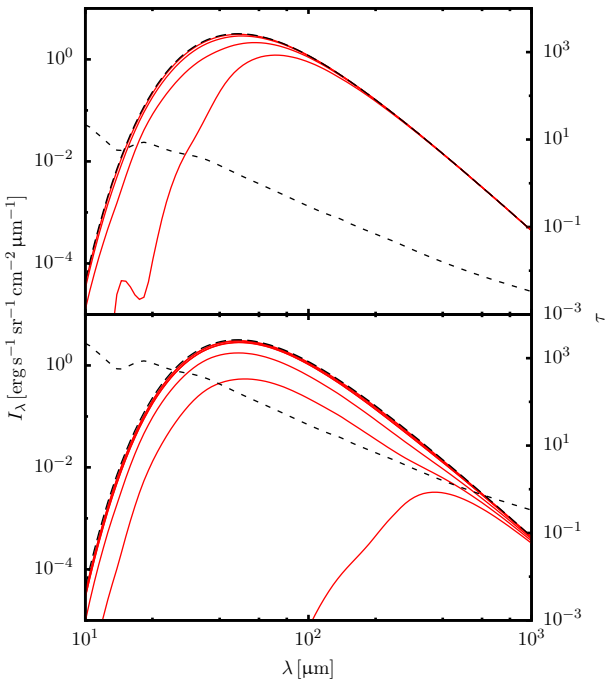

Figure 5. The spectral energy distribution (SED) through the centre of the Bonnor-Ebert spheres, with $M = 1\,\mathrm{M}_\odot$ (top) and $M = 100\,\mathrm{M}_\odot$ (bottom). In both cases, the sphere is illuminated by an undiluted 60 K blackbody. The left y-axis gives the intensity calculated by successive iterations of the algorithm. The lowest red line shows the initial intensity, where each particle has a temperature of 10 K. The lines above show the SED converging towards a 60 K blackbody (black long-dashed line). The right y-axis give the optical depth through the centre of the sphere (short-dashed black line).

4.1.3 Results

Fig. 5 shows the spectral energy distribution (SED) though the centre of the sphere. In the optically thin case, we see that the SED converges on a 60 K blackbody after about three iterations. In the optically thick case, the SED converges after about five iterations.

Fig. 6 shows the particle temperatures after the first three iterations of the algorithm with the optically thin sphere. This includes calculations with both $N_\gamma = 10^6$ and

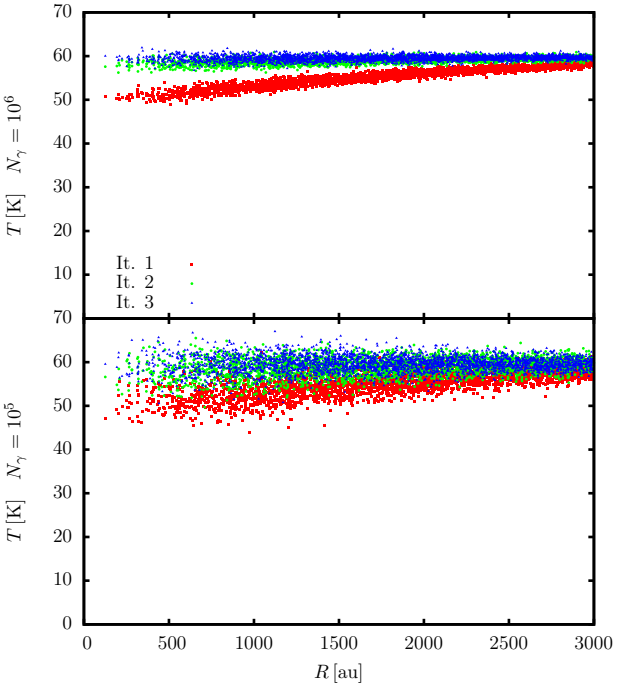

Figure 6. Particle temperature versus radius for optically thin sphere. The top frame shows the temperature calculation with 10^6 luminosity packets. The bottom frame shows the calculation repeated with 10^5 luminosity packets. The particle temperatures from the first iteration are plotted as red squares, from the second iteration as green circles and from the third iteration as blue triangles.

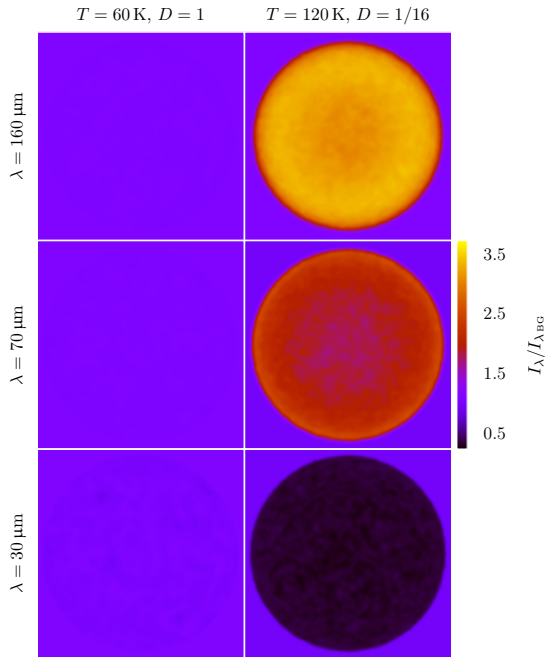
$T = 60\,\mathrm{K}$, $D = 1$
$T = 120\,\mathrm{K}$, $D = 1/16$
$\lambda = 160\,\mu\mathrm{m}$
$\lambda = 70\,\mu\mathrm{m}$
$\lambda = 30\,\mu\mathrm{m}$
$I_\lambda / I_{\lambda\,\mathrm{BG}}$

Figure 7. Intensity maps of a sphere of $N_{\mathrm{PART}} = 3 \times 10^5$ particles with $M = 100\,\mathrm{M}_\odot$ and $R = 2000\,\mathrm{au}$. The radial density profile follows that of a critical Bonnor-Ebert sphere. The left column shows the intensity when the sphere is illuminated by a 60 K blackbody radiation field. The right column shows the intensity when the sphere is illuminated by a 120 K blackbody radiation field, diluted by a factor of 1/16. The top row gives the intensity at $\lambda = 160\,\mu\mathrm{m}$, the middle row gives the intensity at $\lambda = 70\,\mu\mathrm{m}$ and the top row gives the intensity at $\lambda = 30\,\mu\mathrm{m}$. The colour scale gives the intensity, normalised by the background intensity.

$N_\gamma = 10^5$. When $N_\gamma = 10^6$, the particle temperatures settle on a narrow distribution of $T = 59.7 \pm 0.4$K. When $N_\gamma = 10^5$ the temperatures distribution is $T = 59.5 \pm 1.1\,\mathrm{K}$. This demonstrates the Poisson-like uncertainties on Monte Carlo calculations, i.e. a factor of N fewer luminosity packets increases the signal-to-noise ratio by a factor of $\sqrt{N}$.

4.1.4 Diluted blackbody radiation field

We repeat the optically thick simulation, this time with a 120 K blackbody, diluted by a factor of 1/16. In this instance, we expect the dust to reach local radiative equilibrium at some temperature $T < 120\,\mathrm{K}$. Here, when a luminosity packet is absorbed by the dust, it is usually re-emitted at a longer wavelength. The sphere should therefore appear to glow, relative to the background, at long wavelengths and cast a silhouette at short wavelengths.

Fig. 7 shows intensity maps of the optically thick sphere. The left column shows the results of the undiluted blackbody field, the right column shows the results for the diluted field. The sphere in the undiluted field is almost invisible relative to the background intensity at wavelengths between 30 and 160 µm. There is some visible noise at 30 µm, of order ten percent. Here, the optical depth through the sphere is very high ($\tau \sim 100$) and the intensity is very sensitive to temperature fluctuations across the the sphere's outer surface. The sphere in the diluted field behaves as expected; it is brighter than the background at 160 µm and darker than the background at 30 µm.

4.2 Embedded sextuple system

We demonstrate a more realistic application of the algorithm by calculating the dust emission from a protostellar multiple system. High order multiple systems, i.e. $N_\star \geq 3$, are observed amongst mature field stars (e.g. Tokovinin 2008; Eggleton & Tokovinin 2008). These systems are more common amongst pre-Main Sequence stars, embedded in star forming regions such as Taurus and Ophiuchus (e.g. Leinert et al. 1993; Ratzka, Köhler & Leinert 2005; Kraus et al. 2011). Furthermore, similar systems form routinely in simulations of molecular clouds and cores (e.g. Delgado-Donate et al. 2004; Bate 2009; Lomax et al. 2014, 2015).

4.2.1 Set up

We take a snapshot from one of the SPH simulations presented by Lomax et al. (2014)[2]. Here, a $1.3\,\mathrm{M}_\odot$ core collapses

[2] These simulations use initial conditions drawn from distribution functions that reproduce the observed properties of the cores in Ophiuchus. The simulation used here is No. 52, with episodic radiative feedback.

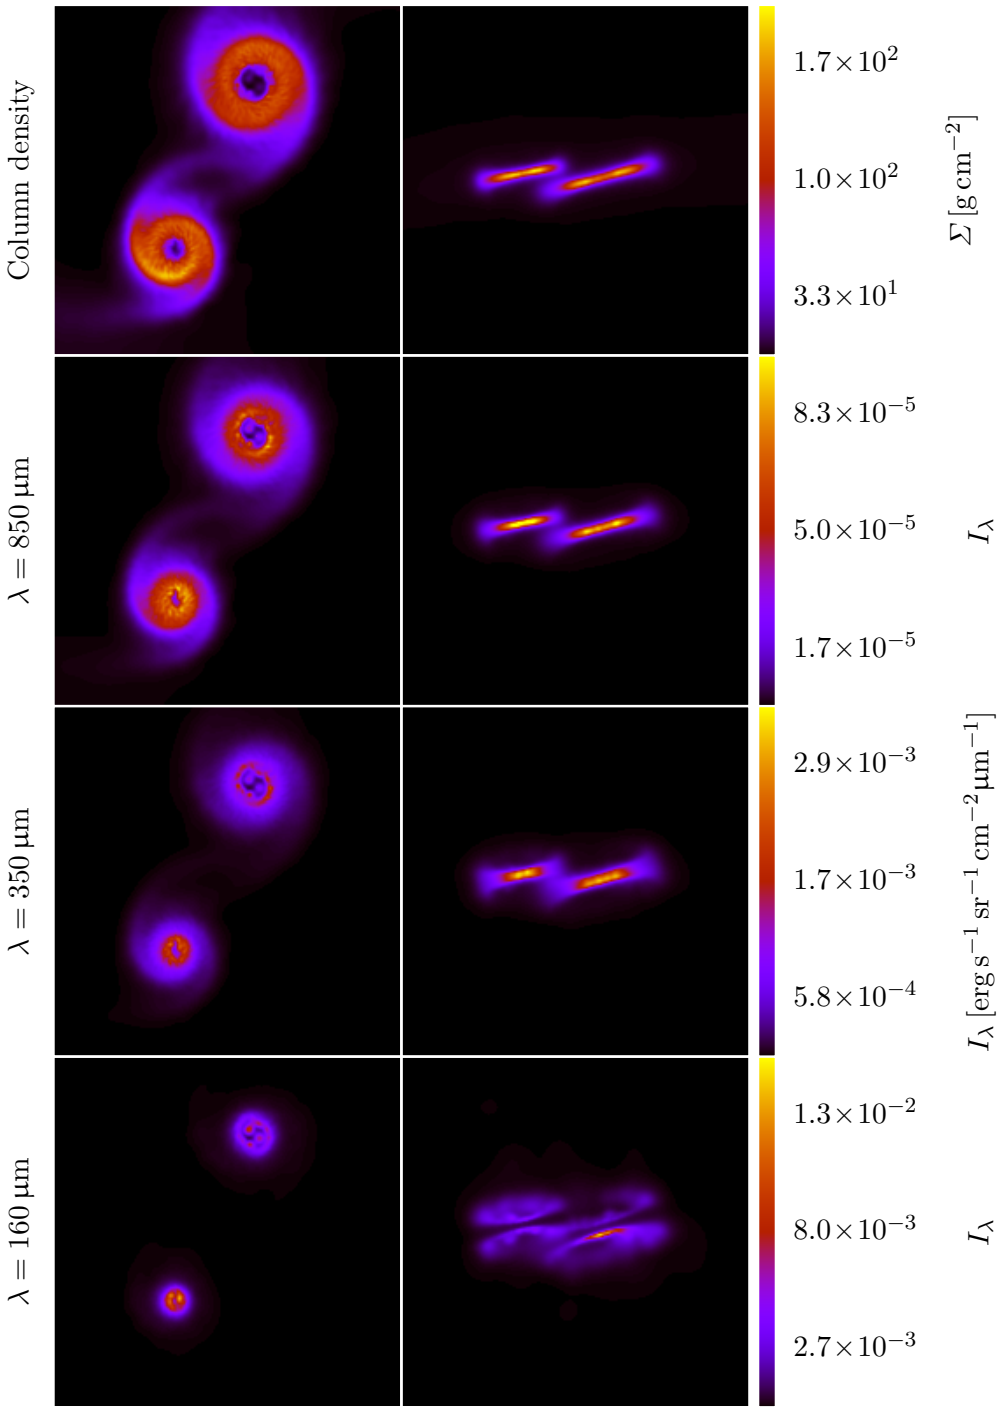

Figure 8. Intensity maps of dust emission from a simulated protostellar sextuple system. The left column shows a face-on view, the right column shows shows an edge-on view. The top row shows the column density of the system. The following three rows show the intensity at $\lambda = 850\,\mu\mathrm{m}$, $350\,\mu\mathrm{m}$ and $160\,\mu\mathrm{m}$. Each frame has area $670\,\mathrm{au} \times 670\,\mathrm{au}$. Note: the colour bar applies to the right column. Figures in the left column have been linearly scaled to fit in the same range.

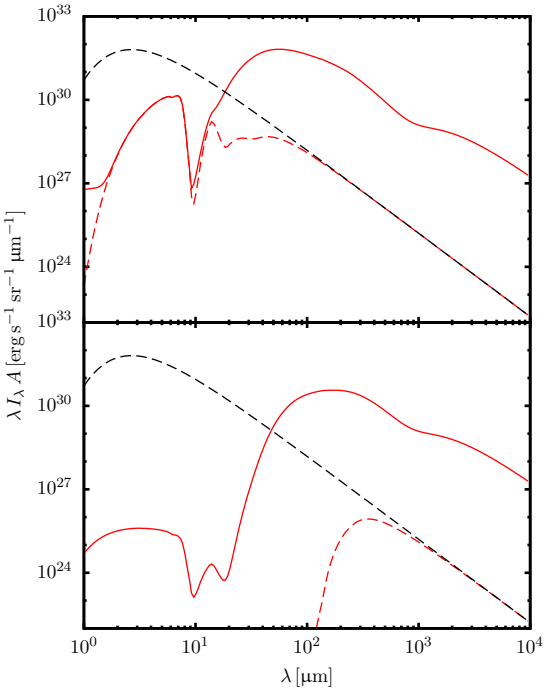

Figure 9. Spectra of dust emission and starlight from the sextuple system. Values are integrated over the map area shown in Fig. 8 . The top frame shows the spectrum from a face-on view. The bottom frame shows the spectrum from an edge-on view. The solid red lines show the total emission. The dashed red line shows the emergent emission from the six protostars. The dashed black line shows the protostellar emission in the absence of dust extinction.

and fragments into seven protostars (represented by sink particles). One of these is ejected from the core and the remaining six form a stable sextuple system. At $t = 10^5$ years (after the initial collapse), the protostars in the sextuple have masses between $0.08\,\mathrm{M}_\odot < M_\star < 0.15\,\mathrm{M}_\odot$. The ejected protostar has mass $M = 0.01\,\mathrm{M}_\odot$. The remainder of the initial mass ($M_{\rm GAS} = 0.65\,\mathrm{M}_\odot$, $N_{\rm PART} = 6.5 \times 10^4$) is distributed between discs and and a diffuse envelope.

We irradiate the system with an undiluted 2.73 K blackbody radiation field and a 10000 K blackbody field, diluted by a factor of 5×10^{-15}. This represents contributions from the Cosmic Microwave Background (CMB) and the galactic stellar population. The protostars are treated as blackbody sources with $R = 4\,\mathrm{R}_\odot$ and $1100\,\mathrm{K} < T < 1500\,\mathrm{K}$; these temperatures are estimated using the Stamatellos, Whitworth & Hubber (2011) episodic accretion model.

4.2.2 Intensity maps and spectra

Fig. 8 shows column density and intensity maps of the embedded system. The system is viewed both face-on and edge-on. In the face-on view, we see a balanced quadruple system (top right) in orbit with a binary system (bottom left), separated by $S \sim 400\,\mathrm{au}$.

In both face-on and edge-on views, the intensity at 850 and 350 μm roughly traces the column density. At 160 μm, face-on, only the the inner regions of the discs are visible. When viewed edge-on, the discs are opaque and significant intensity is only seen above and below their midplanes.

Fig. 9 shows the emission spectrum of the system. When viewed face on, we may define different wavebands, dominated by different sources of radiation. Between $10^3\,\mu\mathrm{m} < \lambda \leq 10^4\,\mu\mathrm{m}$, the dust is optically thin and most of the radiation is from the CMB. Between $10\,\mu\mathrm{m} < \lambda \leq 10^3\,\mu\mathrm{m}$, most of the radiation is from dust emission. Starlight is largely unattenuated in these wavebands, but its fraction of the total emission is very low. At wavelengths less than $10\,\mu$m, the majority of the emission is attenuated starlight. Similar dust emission is seen when the system is viewed edge on. However, the starlight is almost completely extinguished at $\lambda < 200\,\mu$m. We note that the level of extinction at $\lambda < 10\,\mu$m is probably exaggerated as we have not captured scattered light in these spectra. Nevertheless, in the face-on view, we still see the double peaked spectra typical of protostellar discs (e.g. Robitaille 2011; Whitney et al. 2013).

5 FUTURE DEVELOPMENTS

5.1 Performance

While the code is not yet heavily optimized, the algorithm is efficient enough to run on an Intel i5-4200U dual core CPU at 1.60 GHz in a reasonable time. A single iteration of the optically thick sphere in §4.1.1 requires 3.5×10^{-4} s per packet per CPU core. The optically thin sphere requires 5.0×10^{-5} s per packet per CPU core. A single iteration of the sextuple system in §4.2 requires approximately 1.5×10^{-5} s per packet per CPU core. We note that the optimal number of packets N_γ is not obvious. In this paper, we select sensible values through a process of trial and error.

A thorough analysis of the processing time required for this algorithm in different scenarios is beyond the scope of this paper. However, we can estimate the time scaling for simple systems such as a star embedded in a uniform density sphere. Here,

$$t(\bar{\tau}, N_\gamma, \xi, N_{\rm SPH}) \propto \bar{\tau}^2\, N_\gamma\, \xi^2\, N_{\rm SPH}^{1/3}\, \log N_{\rm SPH}. \qquad (5.1)$$

The flight path of a packet is proportional to the square of the average optical depth through the system $\bar{\tau}$. The number of particles along the path is proportional to the square of the kernel extent ξ and the cube root of the total number of particles $N_{\rm SPH}$. The time taken to find each particles scales roughly with the log $N_{\rm SPH}$.

One may ensure that each particle is visited the roughly same number of times, regardless of $N_{\rm SPH}$, by setting $N_\gamma \propto N_{\rm SPH}^{2/3}$. Ignoring ξ, the time scaling is now $t(\tau, N_{\rm SPH}) \propto \bar{\tau}^2\, N_{\rm SPH}\, \log N_{\rm SPH}$. This is similar to that of an SPH simulation timestep. Therefore it may be feasible in the future to run this algorithm on-the-fly in an SPH simulation, especially in circumstances where $\bar{\tau} \lesssim 1$.

5.2 Features

At present, we are able to generate intensity maps of the dust emission, plus attenuated light from background and point sources. However, it is desirable to also capture light

scattered by dust. This can be achieved by modifying Eqn. 3.9. Here, the term $B_\lambda(T)\kappa_\lambda$ represents the amount of dust emission per unit solid angle, per unit wavelength, per unit mass. We may add an extra term which accounts for scattered light:

$$B_\lambda(T)\kappa_\lambda \rightarrow B_\lambda(T)\kappa_\lambda + \frac{1}{\Delta\lambda}\frac{\varepsilon_0}{\Delta t}\frac{1}{m}\sum_{j,\lambda,\Delta\lambda}\Phi(\theta_j)\sigma_{\lambda j}\varsigma_j\,. \quad (5.2)$$

Here, $\sigma_\lambda \equiv a_\lambda\,\chi_\lambda$ is the mass scattering coefficient and θ_j is the angle between luminosity packet j and the viewing angle. The sum is only over luminosity packets with wavelengths in the interval $(\lambda, \lambda + \Delta\lambda)$. The scattering term must be calculated during the Monte Carlo iteration, so the viewing angle and wavelength interval must be known beforehand.

Other future developments include use of the partial diffusion approximation and modified random walk for regions of high optical depth (Min et al. 2009). The SPAMCART code will shortly be made open-source and uploaded to the GitHub repository.

6 SUMMARY

We have developed a method for performing MCRT calculations directly on a distribution of SPH particles The algorithm operates differently from uniform-density cell methods, but the two schemes are mathematically equivalent. This allows an MCRT calculation to be performed on an SPH snapshot with (i) no loss in density resolution and (ii) no introduction of noise from mapping particles to cells.

We present a version of this algorithm that uses the Lucy (1999) method to compute (i) the propagation of luminosity packets through a medium and (ii) the radiative equilibrium temperature. The trajectories of the packets rely on the temperature of the medium so the calculation must be solved by iteration.

We provide two example calculations using the smoothed particle MCRT method. First, we show that a cloud is invisible when it is bathed in an undiluted blackbody radiation field. This holds for both optically thin and optically thick cases. Therefore the code obeys Kirchhoff's law of thermal radiation. We note that this is a powerful test of any radiative transfer code and can be applied to any configuration. Conversely, a cloud which has reached radiative equilibrium with a diluted blackbody field glows at long wavelengths and casts a silhouette at short wavelengths. Second, we generate intensity maps and spectra of protostellar and dust emission from an embedded sextuple system. Here, a double peaked disc spectrum is seen when the system is observed face on. When viewed edge on, the opacity of the disc blocks nearly all of the starlight.

Future additions to the code include the addition of scattered light to synthetic observations and optimisations for optically thick regions of dust. The code, written in Fortran 2003/08 with OpenMP parallelisation, will be made publicly available in the near future.

ACKNOWLEDGEMENTS

OL and APW gratefully acknowledge the support of a consolidated grant (ST/K00926/1) from the UK STFC. We also thank the referee for their considerate and constructive comments.

REFERENCES

Altay G., Croft R. A. C., Pelupessy I., 2008, ACP08, 386, 1931
Barnes J., Hut P., 1986, Nature, 324, 446
Bate M. R., 2009, MNRAS, 392, 590
Delgado-Donate E. J., Clarke C. J., Bate M. R., Hodgkin S. T., 2004, MNRAS, 351, 617
Dobbs C. L., Bonnell I. A., Pringle J. E., 2006, MNRAS, 371, 1663
Dobbs C. L., Burkert A., Pringle J. E., 2011, MNRAS, 417, 1318
Dullemond C. P., 2012, RADMC-3D: A multi-purpose radiative transfer tool. Astrophysics Source Code Library, record ascl:1202.015
Eggleton P. P., Tokovinin A. A., 2008, MNRAS, 389, 869
Ercolano B., Barlow M. J., Storey P. J., Liu X.-W., 2003, MNRAS, 340, 1136
Forgan D., Rice K., 2010, MNRAS, 406, 2549
Gingold R. A., Monaghan J. J., 1977, MNRAS, 181, 375
Harries T. J., 2011, MNRAS, 416, 1500
Henyey L. G., Greenstein J. L., 1941, ApJ, 93, 70
Hubber D. A., Ercolano B., Dale J., 2016, MNRAS, 456, 756
Hubber D. A., Rosotti G. P., 2013, GANDALF (Graphical Astrophysics code for N-body Dynamics And Lagrangian Fluids). GitHub
Kraus A. L., Ireland M. J., Martinache F., Hillenbrand L. A., 2011, ApJ, 731, 8
Leinert C., Zinnecker H., Weitzel N., Christou J., Ridgway S. T., Jameson R., Haas M., Lenzen R., 1993, A&A, 278, 129
Li A., Draine B. T., 2001, ApJ, 554, 778
Lomax O., Whitworth A. P., Hubber D. A., Stamatellos D., Walch S., 2014, MNRAS, 439, 3039
Lomax O., Whitworth A. P., Hubber D. A., Stamatellos D., Walch S., 2015, MNRAS, 447, 1550
Lucy L. B., 1977, ApJ, 82, 1013
Lucy L. B., 1999, A&A, 344, 282
Min M., Dullemond C. P., Dominik C., de Koter A., Hovenier J. W., 2009, A&A, 497, 155
Monaghan J. J., Lattanzio J. C., 1985, A&A, 149, 135
Price D. J., Monaghan J. J., 2004, MNRAS, 348, 139
Ratzka T., Köhler R., Leinert C., 2005, A&A, 437, 611
Robitaille T. P., 2011, A&A, 536, A79
Rundle D., Harries T. J., Acreman D. M., Bate M. R., 2010, MNRAS, 407, 986
Springel V., 2005, MNRAS, 364, 1105
Springel V., Hernquist L., 2003, MNRAS, 339, 312
Stamatellos D., Whitworth A. P., 2005, A&A, 439, 153
Stamatellos D., Whitworth A. P., Hubber D. A., 2011, ApJ, 730, 32
Tokovinin A., 2008, MNRAS, 389, 925
Weingartner J. C., Draine B. T., 2001, ApJ, 548, 296
Whitney B. A., Robitaille T. P., Bjorkman J. E., Dong R., Wolff M. J., Wood K., Honor J., 2013, ApJS, 207, 30
Williams A., Barrus S., Morley R. K., Shirley P., 2005, J. Graphics Tools, 10, 49

$N_{\rm LEAF}$	$N_{\rm OCT}$	$N_{\rm FIXED}$	Levels	$M_{\rm GRID}/M_{\rm SPH}$	$\rm Error_{90\%}$
1	154281	65536^3	16	0.80	0.24
8	22632	8192^3	13	0.87	0.45
64	3242	4096^3	12	0.88	0.76
Gather	–	–	–	–	0.18

Table A1. Statistics of the particle octrees. The first column give the maximum number of particles per leaf. The second column gives the total number of leaf cells. The third column gives the number of cells in an equivalent regular grid. The forth column gives the ratio of the total grid mass to the total particle mass. The fifth column gives the 90th centile density error (see Fig. A2).

APPENDIX A: GRIDDING ERRORS

Transposing an ensemble of SPH particles onto a grid introduces errors to the density field. Here, we perform a brief analysis on these errors. We generate an octree around the distribution of particles used to model the sextuple system in §4.2. First, we build the smallest cuboidal *root* cell that contains all of the particles. This cell is recursively subdivided into eight equal volume cells until each *leaf* cell contains $N_{\rm LEAF}$ particles or fewer. Fig. A1 shows the x-y projection of the particles and the octree with $N_{\rm LEAF} = 8$.

The density of each cell $\rho_{\rm GRID}$ is computed using an SPH *scatter* calculation (see Eqn. 3.4) at the centre of the cell. We compare this with the particle density $\rho_{\rm SPH}$ (scatter calculation) of each particle in the cell. Fig. A2 shows the relative difference between $\rho_{\rm SPH}$ and $\rho_{\rm GRID}$ for the ensemble of particles with different values of $N_{\rm LEAF}$. Note that the particles are sorted into ascending order of density error. For comparison, we also include the relative difference between the particle scatter density and its *gather* density (see Eqn. 3.2)[3]. In Table A1 we give some further statistics including the number of leaf cells, the maximum depth of the tree, the total mass of the grid (defined as the sum of the products of leaf cell volume and leaf cell density) and the 90th centile density error.

Gridding the density field with $N_{\rm LEAF} = 1$, 8 and 64 incurs a 90th centile density error of 24%, 45% and 76% respectively. Furthermore, the mass of the grid is roughly 80% to 90% that of the particle ensemble. We note that although the method presented in the paper is more accurate than gridding an ensemble of particles, MCRT calculations on octrees are mathematically and computationally simpler.

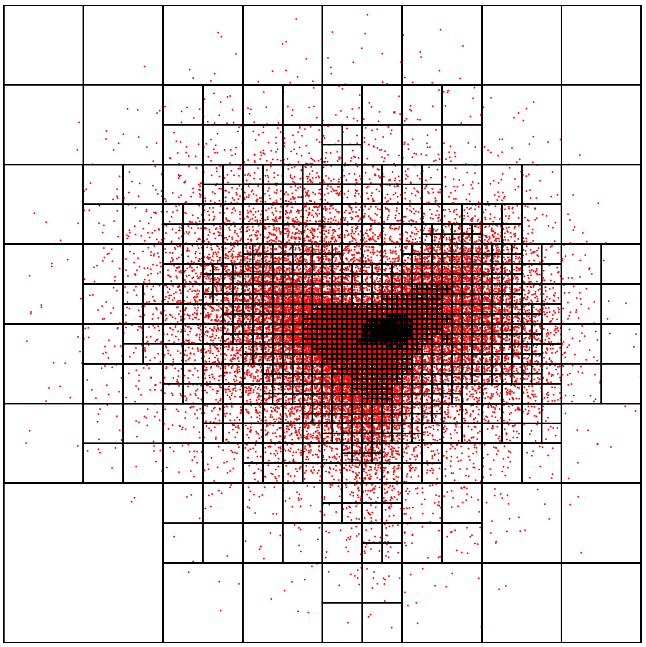

Figure A1. An x-y projection of the octree and particle ensemble with $N_{\rm LEAF} = 8$. The root cell is roughly cubic with an edge length of 0.2 pc. The smallest leaf cells have an edge-length of 2.4×10^{-5} pc.

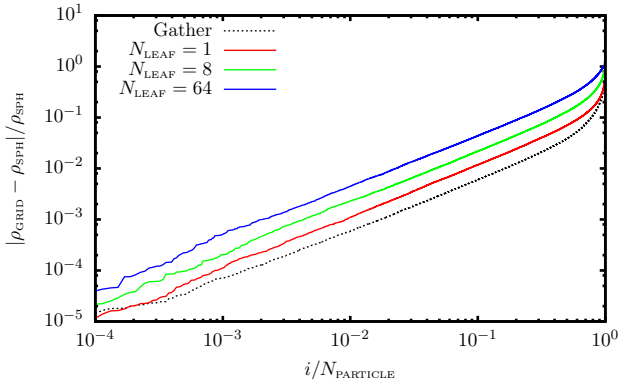

Figure A2. The relative difference between the grid cell density and the SPH particle density against particle number. Particles are sorted in order of ascending error. From top to bottom, the solid coloured lines show $N_{\rm LEAF} = 64$, 8 and 1. The dotted black line shows the relative difference between the gather and scatter density for each particle, i.e. $|\rho_{\rm GATHER} - \rho_{\rm SCATTER}|/\rho_{\rm SCATTER}$.

APPENDIX B: KERNEL FUNCTION

The M4 kernel has support $s \leq 2$, where $s = r/h$. The density at s is given by

$$w(s) = \frac{1}{\pi}\left(\frac{1}{4}\max(2-s,0)^3 - \max(1-s,0)^3\right), \quad \text{(B1)}$$

and the density gradient is given by

$$\frac{\mathrm{d}}{\mathrm{d}s}w(s) = -\frac{1}{\pi}\left(\frac{3}{4}\max(2-s,0)^2 - 3\max(1-s,0)^2\right). \quad \text{(B2)}$$

The column density from dimensionless impact parameter $c = b/h$ to s is given by

$$W(c,s) = \frac{1}{\pi}\begin{cases} W_{\rm INNER}(c,s), & c \leq 1,\ s \leq 1, \\ W_{\rm INNER}(c,1) + W_{\rm OUTER}(c,s), & c \leq 1,\ s > 1, \\ W_{\rm OUTER}(c,s), & c > 1,\ s > 1, \\ 0, & c > 2, \end{cases} \quad \text{(B3)}$$

[3] The difference between the scatter and gather density calculations does not indicate that SPH is somehow wrong. It reflects the difference between the mass to volume ratio of an SPH particle (gather), and the superposition of multiple smoothing kernels at a given point in space (scatter).

where $c \leq s$ and

$$\begin{aligned} W_{\rm INNER}(c,s) &= \sqrt{s^2-c^2} \\ &\times \left(\frac{3}{16}s^3 - \frac{1}{2}s^2 + \frac{9}{32}c^2 s - c^2 - 1\right) \\ &+ \frac{9}{32}c^4 \log\left(\frac{s+\sqrt{s^2-c^2}}{c}\right), \\ W_{\rm OUTER}(c,s) &= \sqrt{s^2-c^2} \\ &\times \left(-\frac{1}{16}s^3 + \frac{1}{2}s^2 - \frac{3}{32}(c^2+16)\,s + c^2 + 2\right) \\ &- \frac{3}{32}(c^2+16)\,c^2 \log\left(\frac{s+\sqrt{s^2-c^2}}{c}\right). \end{aligned} \tag{B4}$$

In practice, Eqn. B3 can either be computed on-the-fly or stored in a triangular lookup table in the the range $0 \leq c^2 \leq 4$ and $c^2 \leq s^2 \leq 4$.

APPENDIX C: K-D TREE

We use a k-d tree (in three dimensions) to find the particles intersected by a ray. We construct the tree by placing the entire ensemble of particles into a *root* cell. The root cell contains two *branch* cells. We select the dimension k across which the particle position $\boldsymbol{x}$ has the greatest variance. All particles with x_k up to and including the median value are placed in the first cell and the remaining particles are placed in the second cell. This process is repeated recursively until each cell contains 8 or fewer particles, in which case they are *leaf* cells. For each cell, we calculate the Axis-Aligned Bounding Box (AABB) that encompasses all of the particle smoothing volumes within the cell.

We find the particles intersected by a ray by using the slab method (e.g. Williams et al. 2005). Consider a ray with origin $\boldsymbol{o}$, direction $\boldsymbol{n}$ and length l. Also, consider an AABB with lower limits $\boldsymbol{b}^{\rm MIN}$ and upper limits $\boldsymbol{b}^{\rm MAX}$. Along each dimension k, we define a slab (a slab is a space between two parallel planes) with lower limit $b_k^{\rm MIN}$ and upper limit $b_k^{\rm MAX}$. The length of the ray segment, $t_k^{\rm MAX}$-$t_k^{\rm MIN}$, within the slab can be calculated:

$$\begin{aligned} t_k^{\rm MIN} &= \min([b_k^{\rm MIN} - o_k]/n_k,\, [b_k^{\rm MAX} - o_k]/n_k)\,, \\ t_k^{\rm MAX} &= \max([b_k^{\rm MIN} - o_k]/n_k,\, [b_k^{\rm MAX} - o_k]/n_k)\,. \end{aligned} \tag{C1}$$

The length of the ray through the AABB in all three dimensions, $T_{\rm MAX} - T_{\rm MIN}$, is given by

$$\begin{aligned} T_{\rm MIN} &= \max(t_1^{\rm MIN}, t_2^{\rm MIN}, t_3^{\rm MIN})\,, \\ T_{\rm MAX} &= \min(t_1^{\rm MAX}, t_2^{\rm MAX}, t_3^{\rm MAX})\,. \end{aligned} \tag{C2}$$

The ray intersects the AABB if the following statements are true:

$$\begin{aligned} T_{\rm MAX} &> T_{\rm MIN}\,, \\ T_{\rm MIN} &< l\,, \\ T_{\rm MAX} &> 0\,. \end{aligned} \tag{C3}$$

Starting at the root cell, we check for an AABB-ray intersection. If the result is true, we open the branch cells and recursively repeat the process until we encounter a leaf cell. If the leaf cell is intersected by the ray, the cell's contents are added to a particle list. This method is not exclusive to k-d trees and may be used with any axis aligned tree, such as an octree (e.g. Barnes & Hut 1986).

APPENDIX D: ROOT-FINDING METHOD

Newton's method can be used to find the length of a luminosity packet trajectory. The equation $\Sigma(l) = \tau\,\chi_\lambda^{-1}$, where $\Sigma(l)$ is the column density, can be solved rapidly by iterating

$$l_{n+1} = l_n - \frac{\Sigma(l_n) - \tau\,\chi_\lambda^{-1}}{\Sigma'(l_n)}\,, \tag{D1}$$

where $\Sigma'(l_n) = \rho(\boldsymbol{o} + l_n\boldsymbol{n})$. This performs well if l_n is near the root of $\Sigma(l) - \tau\,\chi_\lambda^{-1}$, but may fail to converge for poor initial choices of l_n.

We modify this method so that convergence is guaranteed. First we bracket the root, i.e. find values l_{n-1} and l_n so that $\Sigma(l_n)$ and $\Sigma(l_{n-1})$ are on opposite sides of $\tau\,\chi_\lambda^{-1}$. We assume that Eqn. 3.7 succeeds in overestimating l and set $l_{n-1} = 0$ and $l_n = l_0$. We now define the parabolic curve which (i) passes through point $[l_n, \Sigma(l_n)]$ with gradient $\Sigma'(l_n)$ and (ii) passes through point $[l_{n-1}, \Sigma(l_{n-1})]$. This curve is described by the equation

$$\begin{aligned} &\Sigma(l_n) - \Sigma(l_{n-1}) \\ &\quad + \Sigma'(l_n)(l_{n-1} - l_n) \\ &\quad + \frac{1}{2}\Sigma''(l_n)(l_{n-1} - l_n)^2 = 0\,, \end{aligned} \tag{D2}$$

and by construction must have a single root in the interval (l_{n-1}, l_n). Rearranging for the second derivative,

$$\Sigma''(l_n) = 2\left[\frac{\Sigma(l_{n-1}) - \Sigma(l_n)}{(l_{n-1} - l_n)^2} - \frac{\Sigma'(l_n)}{l_{n-1} - l_n}\right]. \tag{D3}$$

Solving the quadratic formula, the iteration function is now

$$l_{n+1} = l_n - \frac{2\,[\Sigma(l_n) - \tau\,\chi_\lambda^{-1}]}{\Sigma'(l_n) + \sqrt{[\Sigma'(l_n)]^2 - 2\,[\Sigma(l_n) - \tau\,\chi_\lambda^{-1}]\,\Sigma''(l_n)}}\,. \tag{D4}$$

If $\Sigma(l_n)$ and $\Sigma(l_{n+1})$ are on opposite sides of $\tau\,\chi_\lambda^{-1}$, the new l_{n-1} is set to the old l_n. Otherwise l_{n-1} remains the same.